~~STREAMING~~ MEMORY SYSTEM

BACKGROUND OF THE INVENTION

Priority is claimed under 35 U.S.C. 119(e) based on provisional application no. 60/100,147 filed September 14, 1998.

1. Field of the Invention

The present invention relates to a memory control and, more particularly, a memory control system that allows faster access to memory cells.

2. Description of Related Art

Modern computer systems use well known dynamic memory chips that are arranged as a matrix of rows and columns. FIGS. 1-3 illustrates such dynamic memory chips. As illustrated in FIG. 1, the various memory cells are disbursed within different banks, such as banks 1-8 illustrated in FIG. 1. FIG. 2 illustrates certain of the components within a given bank of the memory. As is known, associated with the bank will be a row decoder 22, a column decoder 24, a memory array 26, a **sense** amplifier 28, and column selection circuitry 30, also referred to herein as I/O circuitry.

A portion of the memory array 26 is further illustrated in FIG. 3, as shown, a plurality of memory cells 30 are attached to particular rows 32, (also known as word lines) and columns 34 (also known as bitlines).

Operation of memories as described above is well known, for example that an address will be input to the memory array, a memory cell associated with the input address will be accessed, and the data stored in that memory cell can be read out. Similarly, if data needs to be written into the memory array, the data will have an associated address, and that address will be used to store the data into the memory cell associated with that particular address.

Memory chips as described above also use techniques in order to increase their speed and efficiency. Using conventional techniques, such chips can access a word of data in a different column of a pre-selected row in a very efficient manner (typically one word per cycle) but access to a word in a different (non-selected) row is relatively slow (typically ten cycles). Furthermore, since these chips are divided into banks as mentioned previously, which each include separate row/column matrices, as illustrated in FIG. 3, this allows for a row access to be performed on one bank while column accesses are being made to a different bank. While such operation improves efficiency somewhat, improvements are still needed.

Particularly, conventional memory systems process memory operations in the order they are received. If they receive addresses in the same row as the previous address, then a column access is performed. Otherwise, a row access is performed. While this mode of operation yields adequate performance for access patterns with significant spacial locality, performances are degraded by almost 90% for unstructured address streams. Such unstructured address streams are typically of indirect vector or stream references.

Accordingly, an improved streaming memory system is needed that allows for improved performance in accessing unstructured address streams.

BRIEF SUMMARY OF THE INVENTION

It is an object of the present invention to increase memory performance for access to unstructured address streams.

Another object of the present invention to complete memory access operations in an order different from the order that such memory access operations were requested.

In order to achieve the above object of the inventions, among others, the present invention is a memory system that receives addresses corresponding to data in an order. The memory system includes an address buffer that receives addresses in the order provided by a computer system, a memory array, a control circuit that presents addresses to the memory array in an order different than the order in which they were received by the address buffer; and a read buffer that receives data read out from the memory array.

BRIEF DESCRIPTION OF THE DRAWINGS

FIGS. 1-3 illustrate a conventional dynamic random access memory device.

FIG. 4 illustrates a streaming memory system according to the present invention.

DETAILED DESCRIPTION OF PRESENTLY PREFERRED EMBODIMENTS

FIG. 4 illustrates the streaming memory system according to the present invention. The invention includes an address buffer 52, memory 54, a read buffer 56, and

a schedule/control circuit 58, each of which are described in further detail hereinafter. While in the presently preferred embodiment, the memory 54 is a discrete semiconductor circuit chip, it is contemplated that the streaming memory system according to the present invention can also be implemented on a single integrated chip.

With reference to FIG. 4, the address buffer 52 inputs addresses that will be used by the memory 54 to access certain particular memory cells during a read or write operation. The address buffer 52 must be large enough to store addresses for a predetermined number of memory accesses that have been provided by the overall computer system 60, such as four, eight or more.

Whereas in conventional systems addresses are supplied to the memory device and operated upon in the supplied order, the present invention includes an address buffer 52 which receive and queues a number of addresses, such as 8 different addresses. The present invention has the capability of advantageously operating upon these addresses in an order different from that received by the address buffer 52, which helps promote efficient operation as has been mentioned and will be described further hereinafter. Because the addresses will be operated upon in an order different than that received, the present invention also includes a read buffer 56, which stores the data that is read out from the memory device 54. This allows, therefore, the data that is read out to be stored and subsequently transmitted from the read buffer 56 to the computer system 60 in the same order as the order in which the addresses were received by the address buffer 52 from the computer system 60, for those addresses corresponding to read operations. The schedule/control circuit 58 controls the operation of re-ordering the addressing, as will now be further described. It should be noted, however, that the operations that are

conventionally required in order to access a memory device are not described in detail. In fact, such conventional operations will differ depending upon the type of memory device that exists, such as conventional dynamic random access memories, synchronous dynamic random access memories, or rambus dynamic random access memories, as well as other types of memories that are accessed as one or more banks of rows and columns.

On initiation of a new cycle, a new address is input to the address buffer 52 and the comparison of addresses in the address buffer 52 with the address of the active row in the previous cycle is made using a comparator that compares the row address of each address with the row address of the currently active row. If one of more of the addresses in the address buffer 52 correspond to an address associated with the active row from the previous cycle (also termed currently active row), the schedule/control circuit 58 will initiate those control signals required to perform column addressing of the oldest address in the address buffer 52 that corresponds to the active row using a priority encoder that selects the first address entered into the address buffer that is contained in the currently active row. That address, therefore, will be operated upon during that cycle be input into memory 54 so that the memory cell associated with that address can be accessed. Thus, the comparator and priority encoder, which are implemented preferably using hardwired logic that make up the schedule/control circuit 58, operate every cycle in parallel to select an access to be run. Each cycle the logic scans the addresses in the address buffer 52 to find one (if any) that is to an active row and selects this address for a column access.

If a simultaneous row access is also possible, the logic in the schedule/control circuit 58 also scans the addresses to find one for which a row access

would be profitable, one for which there are no more addresses to the active row in its bank and (optionally) for which there are several other addresses in the same row queued. Thus, for certain memory devices, while column addressing of an address is performed, the schedule/control circuit 58 can also initiate row addressing of a row in a bank other than the currently active row bank. For instance, if column addressing of an address associated with the active row is accessing data in bank 1, the schedule/control circuit 58 may initiate row addressing for a row within bank 6, since bank 6 does not currently contain an active row, so that in a subsequent cycle column addressing of that row can take place. Thus, according to the present invention, row access latency can be hidden under column accesses to other banks, thus improving efficiency of this system.

Once a row access is initiated, the schedule/control circuit 58 will also initiate subsequent control operations, depending on whether a read or a write operation was to take place.

If a write operation takes place, the associated data is written into the addressed memory cell location.

If a read operation takes place, the schedule/control circuit will cause the read out of data from the addressed memory cell location, and storage of that data into a read buffer 56 so that the data can read out of the read buffer 56 in the order that the read addresses were initially received into the address buffer 52. More particularly, in a read operation, the read buffer 56 is indexed by a pair of pointers in a manner such as that used to reorder instructions in processors that allow out of order execution. As each read access is inserted into the address buffer 52, the next sequential location in the read buffer 56 is identified by a read-tail pointer, reserved for this access, and marked pending. The

value of the read-tail pointer is queued with the address in the address buffer 52 to record the location assigned and the read-tail pointer is incremented modulo the size of the read buffer 56. When the queued read access is actually performed, the data read is inserted into the read buffer 56 location reserved for this access using the pointer queued with the address in the address buffer 52 and this location is marked full.

A read-head pointer is used to remove data from the read buffer 56. On reset the read-head and read-tail pointer both point to the same location and that location is marked empty. As read accesses arrive, the read-tail pointer is incremented by the schedule/control circuit 58 and locations are marked pending to allocate sequential read buffer 56 locations to these sequential accesses. Finally, as read accesses are performed, some of these pending locations are filled.

Whenever the location identified by the read-head pointer is marked full, the value in that location is output, the location marked empty, and the read-head pointer incremented modulo the size of the read buffer 56. Because the read-data for the accesses is output in the same order that the read addresses arrived, ordering, read order is preserved even though memory accesses are performed out of order. The read buffer 56 in effect reorders the out of order memory accesses.

Ordering is also important in a memory system, and the present invention allows for completion of memory write operations out of order with operations that read or write other memory rows because these operations are guaranteed not to be of the same address. The schedule/control circuit 58 always performs accesses in the same row in the original requested order, thus preserving the original order for two writes to the same location or a read and a write to a given location. The present system preserves the

ordering of write operations and the relative ordering of reads and writes by always scanning for accesses to an active row in the order that accesses arrived. Thus, an 'older' access to a given row, and hence a given location, will always occur before a later access to the same row, and hence same location. Only accesses to distinct rows and hence distinct locations are reordered. Thus, read before write or write before read hazards are not a problem with the schedule/control circuit 58 re-ordering.

In a preferred implementation of the present invention, the address buffer and the read buffer are partitioned so that addresses and data for each of the separate banks are buffered separately. Thus, the schedule/control circuit can access the partitioned buffer associated with a currently active bank, perform an equality comparison on a those address bits necessary to determine if another address in the partitioned buffer corresponds to the currently active row in the currently active bank in order to determine whether to perform fast column addressing for the currently active bank using the another address, or instead initiate addressing of another bank in the manner previously described such that row addressing of another bank is performed while the fast column addressing of the bank that previously had an associated active row is ongoing.

The present invention also contemplates initiating access of multiple rows in different banks at the same time.

Although the present invention has been described in detail with reference to the specific embodiments and examples provided herein, those skilled in the art will realize that various modifications and/or substitutions could be made to such specific

examples while remaining within the spirit and scope of the invention as set forth in the appended claims.